

# A System-Level Game Semantics

Dan R. Ghica<sup>1</sup> and Nikos Tzevelekos<sup>2</sup>

<sup>1</sup> University of Birmingham

<sup>2</sup> Queen Mary, University of London

**Abstract.** Game semantics is a trace-like denotational semantics for programming languages where the notion of legal observable behaviour of a term is defined combinatorially, by means of rules of a game between the term (the *Proponent*) and its context (the *Opponent*). In general, the richer the computational features a language has the less constrained the rules of the semantic game. In this paper we consider the consequences of taking this relaxation of rules to the limit, by granting the Opponent *omnipotence*, that is, permission to play any move without combinatorial restrictions. However, we impose an epistemic restriction by not granting Opponent *omniscience*, so that Proponent can have undisclosed secret moves. We introduce a basic C-like programming language and we define such a semantic model for it. We argue that the resulting semantics is an appealingly simple combination of operational and game semantics and we show how certain traces explain system-level attacks, i.e. plausible attacks that are realisable outside of the programming language itself. We also show how allowing Proponent to have secrets ensures that some desirable equivalences in the programming language are preserved.

## 1 Introduction

Game semantics came to prominence by solving the long-standing open problem of *full abstraction* for PCF [1,6] and it consolidated its status as a successful approach to modelling programming languages by being used in the definition of numerous other fully abstract programming language models. The approach of game semantics is to model computation as a formal interaction, called a *game*, between a term and its context. Thus, a semantic game features two players: a *Proponent*, representing the term, and an *Opponent* (O), representing the context. The interaction is formally described by sequences of game moves, called *plays*, and a term is modelled by a corresponding *strategy*, that is, the set of all its possible plays. To define a game semantics one needs to define what are the rules of the game and what are the abilities of the players.

For PCF games, the rules are particularly neat, corresponding to the so-called “principles of polite conversation”: moves are divided into *questions* and *answers*; players must take turns; no question can be asked unless it is made possible (*enabled*) by an earlier relevant question; no answer can be given unless it is to the most recent unanswered question. The legality constraints for plays can be imposed as combinatorial conditions on sequences of moves.

Strategies also have combinatorial conditions which characterise the players rather than the game. They are uniformity conditions which stipulate that if in certain plays P makes a certain move, in other plays it will make an analogous move. The simplest condition is *determinism*, which stipulates that in any strategy if two plays are equal up to a certain P move, their subsequent P moves must also be the same. Relaxing some of the combinatorial constraints on plays and strategies elegantly leads from models of PCF to models of more expressive programming languages. For example, relaxing a condition called *innocence* leads to models of programming language with state [2], relaxing *bracketing* leads to models of programming languages with control [9], and in the absence of *alternation* we obtain languages for concurrency [5].

*Contribution.* In this paper we consider the natural question of what happens if in a game semantics we remove combinatorial constraints from O's behaviour. Unlike conventional game models, our construction is asymmetric: P behaves in a way determined by the programming language and its inherent limitations, whereas O may represent plausible behaviour which is, however, not syntactically realizable neither in the language nor in some obvious extensions. In this paper we will see that such a model is, in a technical sense, well formed and that the notion of equivalence it induces is interesting and useful.

We study such a relaxed game model using an idealized type-free C-like language. The notion of *available move* is modelled using a notion of *secret* similar to that used in models of security protocols, formally represented using *names*. This leads to a notion of Opponent which is omnipotent but not omniscient: it can make any available move in any order, but some moves can be hidden from it. This is akin to Dolev-Yao attacker model of security.

We show how inequivalences in this semantic model capture *system-level attacks*, i.e. behaviours of the ambient system which, although not realizable in the language itself, can be nevertheless enacted by a powerful enough system. Despite the existence of such a powerful ambient system we note that many interesting equivalences still hold. This provides evidence that questions of semantic equivalence can be formulated outside the conventional framework of a *syntactic* context.

Technically, the model is expressed in an operationalised version of game semantics like Laird's [11] and names are handled using nominal sets [4].

## 2 A system-level semantics

### 2.1 Syntax and operational semantics

We introduce a simple untyped C-like language which is just expressive enough to illustrate the basic concepts. A *program* is a list of *modules*, corresponding roughly to files in C. A *module* is a list of function or variable declarations. An exported variable or function name is globally visible, otherwise its scope is the

module. In extended BNF-like notation we write:

$$\begin{aligned} Prog &::= Mod^* & Mod &::= Hdr Mod & Hdr &::= \text{export } \bar{x}; \text{import } \bar{x}; \\ Dcl &::= \text{decl } x = n; Dcl \mid \text{decl } Func; Dcl \mid \epsilon \end{aligned}$$

The header  $Hdr$  is a list of names exported and imported by the program, with  $x$  an identifier (or list of identifiers  $\bar{x}$ ) taken from an infinite set  $\mathcal{N}$  of *names*, and  $n \in \mathbb{Z}$ .

As in C, functions are available only in global scope and in uncurried form:

$$Func ::= x(\bar{x})\{\text{local } \bar{x}; Stm \text{ return } Exp;\}$$

A function has a name and a list of arguments. In the body of the function we have a list of local variable declarations followed by a list of statements terminated by a return statement. Statements and expressions are (with  $n \in \mathbb{Z}$ ).

$$\begin{aligned} Stm &::= \epsilon \mid \text{if}(Exp)\text{then}\{Stm\}\text{else}\{Stm\}; Stm \mid Exp = Exp; Stm \mid Exp(Exp^*); Stm \\ Exp &::= Exp \star Exp \mid *Exp \mid Exp(Exp^*) \mid \text{new}() \mid n \end{aligned}$$

Statements are branching, assignment and function call. For simplicity, iteration is not included as we allow recursive calls. Expressions are arithmetic and logical operators, variable dereferencing ( $*$ ), variable allocation and integer constants. Function call can be either an expression or a statement. Because the language is type-free the distinction between statement and expression is arbitrary and only used for convenience.

If  $\text{decl } f(\bar{x})\{e\}$  is a declaration in module  $M$  we define  $f @ M = e[\bar{x}]$ , interpreted as “the definition of  $f$  in  $M$  is  $e$ , with arguments  $\bar{x}$ .”

A *frame* is given by the grammar below, with  $op \in \{=, \star, ;\}$ ,  $op' \in \{*, -\}$ .

$$t ::= \text{if } (\square) \text{ then } \{e\} \text{ else } \{e\} \mid \square op e \mid v op \square \mid op' \square \mid \square e \mid v \square \mid (\square, e) \mid (v, \square)$$

We denote the “hole” of the frame by  $\square$ . We denote by  $\mathcal{F}s$  the set of lists of frames, the *frame stacks*. By  $v$  we denote *values*, defined below.

Our semantic setting is that of nominal sets [4], constructed over the multi-sorted set of names  $\mathcal{N} = \mathcal{N}_\lambda \uplus \mathcal{N}_\phi \uplus \mathcal{N}_\kappa$  where each of the three components is a countably infinite set of *location names*, *function names* and *function continuation names* respectively. We range over names by  $a, b$ , etc. Specifically for function names we may use  $f$ , etc.; and for continuation names  $k$ , etc. For each set of names  $\mathcal{X}$  we write  $\lambda(\mathcal{X})$ ,  $\phi(\mathcal{X})$  and  $\kappa(\mathcal{X})$  for its restriction to location, function and continuation names respectively. We write  $\nu(x)$  for the *support* of  $x$ , for any element  $x$  of a nominal set  $X$ , i.e. all the free names occurring in it.

A store is defined as a pair of partial functions with finite domain:

$$s \in Sto = (\mathcal{N}_\lambda \multimap_{\text{fn}} (\mathbb{Z} + \mathcal{N}_\lambda + \mathcal{N}_\phi)) \times (\mathcal{N}_\kappa \multimap_{\text{fn}} \mathcal{F}s \times \mathcal{N}_\kappa)$$

The first component of the store assigns integer values (data), other locations (pointers) or function names (pointer to functions) to locations. The second stores *continuations*, used by the system to resume a suspended function call.

We write  $\lambda(s)$ ,  $\kappa(s)$  for the two projections of a store  $s$ . By abuse of notation, we may write  $s(a)$  instead of  $\lambda(s)(a)$  or  $\kappa(s)(a)$ . Since names are sorted, this is unambiguous. The support  $\nu(s)$  of  $s$  is the set of names appearing in its domain or value set. For all stores  $s, s'$  and set of names  $\mathcal{X}$ , we use the notations:

**restrict-to:** only consider the subset of a store defined at a given set of names,

$$s \upharpoonright \mathcal{X} = \{(a, y) \in s \mid a \in \mathcal{X}\}$$

**restrict-from:** only consider the subset of a store that is not defined at a given set of names,  $s \setminus \mathcal{X} = s \upharpoonright (\text{dom}(s) \setminus \mathcal{X})$

**update:** change the values in a store,  $s[a \mapsto x] = \{(a, x)\} \cup (s \setminus \{a\})$  and, more generally,  $s[s'] = s' \cup (s \setminus \text{dom}(s'))$

**valid extension:**  $s \sqsubseteq s'$  if  $\text{dom}(s) \subseteq \text{dom}(s')$

**closure:**  $Cl(s, \mathcal{X})$  is the least set of names containing  $\mathcal{X}$  and all names reachable from  $\mathcal{X}$  through  $s$  in a transitively closed manner, i.e.  $\mathcal{X} \subseteq Cl(s, \mathcal{X})$  and if  $(a, y) \in s$  with  $a \in Cl(s, \mathcal{X})$  then  $\nu(y) \in Cl(s, \mathcal{X})$ .

We give a semantics for the language using a frame-stack abstract machine. It is convenient to take *identifiers* to be *names*, as it gives a simple way to handle pointers to functions in a way much like that of the C language. We define a *value* to be a name, an integer, or a tuple of values:  $v ::= () \mid a \mid n \mid (v, v)$ . The value  $()$  is the unit for the tuple operation. Tupling is associative and for simplicity we identify tuples up to associative isomorphism, so  $(v, (v, v)) = ((v, v), v) = (v, v, v)$  and  $(v, ()) = v$ , etc. If a term is not a value we write it as  $e$ .

The *Program configurations* of the abstract machine are

$$\langle N \mid P \vdash s, t, e, k \rangle \in \mathcal{N} \times \mathcal{N} \times \text{Sto} \times \mathcal{F}s \times \text{Exp} \times \mathcal{N}_\kappa$$

$N$  is a set of *used names*;  $P \subseteq N$  is the set of *public names*;  $s$  is the *program state*;  $t$  is a list of frames called the *frame stack*;  $e$  is the (closed) expression, being evaluated; and  $k$  is a *continuation name*, which for now will stay unchanged.

The transitions of the abstract machine

$$\langle N \mid P \vdash s, t, e, k \rangle \longrightarrow \langle N' \mid P' \vdash s', t', e', k \rangle$$

are defined by case analysis on the structure of  $e$  then  $t$  in a standard fashion, as in Fig. 1. Branching is as in C, identifying non-zero values with true and zero with false. Binary operators are evaluated left-to-right, also as in C. Arithmetic and logic operators ( $\star$ ) have the obvious evaluation. Dereferencing is given the usual evaluation, with a note that in order for the rule to apply it is implied that  $v$  is a location and  $s(v)$  is defined. Local-variable allocation extends the domain of  $s$  with a fresh secret name. Local variables are created fresh, locally for the scope of a function body. The **new**() operator allocates a secret and fresh location name, initialises it to zero and returns its location. The return statement is used as a syntactic marker for an end of function but it has no semantic role.

Structural rules, such as function application and tuples are as usual in call-by-value languages, i.e. left-to-right. Function call also has a standard evaluation. The body of the function replaces the function call and its formal arguments  $\bar{x}$  are substituted by the tuple of arguments  $v'$  in point-wise fashion. Finally, non-canonical forms also have standard left-to-right evaluations.

Case  $e = v$  is a value.

$$\begin{aligned}
\langle N \mid P \vdash s, t \circ (\text{if } (\square) \text{ then } \{e_1\} \text{ else } \{e_2\}), v, k \rangle &\longrightarrow \langle N \mid P \vdash s, t, e_1, k \rangle, \text{ if } v \in \mathbb{Z} \setminus \{0\} \\
\langle N \mid P \vdash s, t \circ (\text{if } (\square) \text{ then } \{e_1\} \text{ else } \{e_2\}), v, k \rangle &\longrightarrow \langle N \mid P \vdash s, t, e_2, k \rangle, \text{ if } v = 0 \\
\langle N \mid P \vdash s, t \circ (\square \text{ op } e), v, k \rangle &\longrightarrow \langle N \mid P \vdash s, t \circ (v \text{ op } \square), e, k \rangle \text{ for } \text{op} \in \{=, *, ;\} \\
\langle N \mid P \vdash s, t \circ (v \star \square), v', k \rangle &\longrightarrow \langle N \mid P \vdash s, t, v'', k \rangle, \text{ and } v'' = v \star v' \\
\langle N \mid P \vdash s, t \circ (v; \square), v', k \rangle &\longrightarrow \langle N \mid P \vdash s, t, v', k \rangle \\
\langle N \mid P \vdash s, t \circ (a = \square), v, k \rangle &\longrightarrow \langle N \mid P \vdash s[a \mapsto v], t, (), k \rangle \\
\langle N \mid P \vdash s, t \circ (*\square), v, k \rangle &\longrightarrow \langle N \mid P \vdash s, t, s(v), k \rangle \\
\langle N \mid P \vdash s, t \circ (\square; e), \text{local } x, k \rangle &\longrightarrow \langle N \cup \{a\} \mid P \vdash s[a \mapsto 0], t, e[a/x], k \rangle, \text{ if } a \notin N \\
\langle N \mid P \vdash s, t \circ (\square(e)), v, k \rangle &\longrightarrow \langle N \mid P \vdash s, t \circ (v(\square)), e, k \rangle \\
\langle N \mid P \vdash s, t \circ ((\square, e)), v, k \rangle &\longrightarrow \langle N \mid P \vdash s, t \circ ((v, \square)), e, k \rangle \\
\langle N \mid P \vdash s, t \circ ((v, \square)), v', k \rangle &\longrightarrow \langle N \mid P \vdash s, t, (v, v'), k \rangle \\
\langle N \mid P \vdash s, t \circ (f(\square)), v', k \rangle &\longrightarrow \langle N \mid P \vdash s, t, e[v'/\overline{x}], k \rangle, \text{ if } f @ M = e[\overline{x}] \quad (F)
\end{aligned}$$

Case  $e$  is not a canonical form.

$$\begin{aligned}
\langle N \mid P \vdash s, t, \text{if } (e) \text{ then } \{e_1\} \text{ else } \{e_2\}, k \rangle &\longrightarrow \langle N \mid P \vdash s, t \circ (\text{if } (\square) \text{ then } \{e_1\} \text{ else } \{e_2\}), e, k \rangle \\
\langle N \mid P \vdash s, t, e \text{ op } e', k \rangle &\longrightarrow \langle N \mid P \vdash s, t \circ (\square \text{ op } e'), e, k \rangle, \text{ if } \text{op} \in \{=, *, ;\} \\
\langle N \mid P \vdash s, t, \text{op } e, k \rangle &\longrightarrow \langle N \mid P \vdash s, t \circ (\text{op } \square), e, k \rangle, \text{ if } \text{op} \in \{\text{return}(-), *\} \\
\langle N \mid P \vdash s, t, \text{new}(), k \rangle &\longrightarrow \langle N \cup \{a\} \mid P \vdash s[a \mapsto 0], t, a, k \rangle, \text{ if } a \in \mathcal{N}_\lambda \setminus N \\
\langle N \mid P \vdash s, t, e(e'), k \rangle &\longrightarrow \langle N \mid P \vdash s, t \circ (\square(e')), e, k \rangle \\
\langle N \mid P \vdash s, t, (e, e'), k \rangle &\longrightarrow \langle N \mid P \vdash s, t \circ ((\square, e')), e, k \rangle
\end{aligned}$$

**Fig. 1.** Operational semantics

## 2.2 System semantics

The conventional function-call rule (F) is only applicable if there is a function definition in the module. If the name used for the call is not the name of a known function then the normal operational semantics rules no longer apply. We let function calls and returns, when the function is not locally defined, be a mechanism for interaction between the program and the ambient system. A *System configuration* is a triple  $\langle\langle N \mid P \vdash s \rangle\rangle \in \mathcal{N} \times \mathcal{N} \times \text{Sto}$ .

Given a module  $M$  we will write as  $\llbracket M \rrbracket$  a transition system defining its system-level semantics (SLS). Its states are  $\mathcal{S}\llbracket M \rrbracket = \text{Sys}\llbracket M \rrbracket \cup \text{Prog}\llbracket M \rrbracket$ , where  $\text{Prog}\llbracket M \rrbracket$  is the set of abstract-machine configurations of the previous section and  $\text{Sys}\llbracket M \rrbracket$  is the set of system configurations defined above. The SLS is defined at the level of modules, that is programs with missing functions, similarly to what is often deemed a *compilation unit* in most programming languages.

Let  $\mathcal{L}_{PS} \simeq \mathcal{L}_{SP} = \{\text{call } f, v, k \mid f \in \mathcal{N}_\lambda, k \in \mathcal{N}_\kappa, v \text{ a value}\} \cup \{\text{ret } v, k \mid k \in \mathcal{N}_\kappa, v \text{ a value}\}$ . The transition relation is of the form:

$$\delta[M] \subseteq (Prog[M] \times Prog[M]) \cup (Prog[M] \times \mathcal{L}_{PS} \times Sto \times Sys[M]) \\ \cup (Sys[M] \times \mathcal{L}_{SP} \times Sto \times Prog[M])$$

In transferring control between Program and System the continuation pointers ensure that upon return the right execution context can be recovered. We impose several hygiene conditions on how continuations are used, as follows. We distinguish between P-continuation names and S-continuation names. The former are created by the Program and stored for subsequent use, when a function returns. The latter are created by the System and are not stored. The reason for this distinction is both technical and intuitive. Technically it will simplify proving that composition is well-defined. Intuitively, mixing S and P continuations does not create any interesting behaviour. If S gives P a continuation it does not know then P can only crash. It is not interesting for S to make P crash, because S can crash directly if so it chooses. So this is only meant to remove trivial behaviour.

The first new rule, called program-to-system call is:

*Program-to-System call:*

$$\langle N \mid P \vdash s, t \circ (f(\square)), v, k \rangle \xrightarrow[s \upharpoonright \lambda(P')]{\text{call } f, v, k'} \langle\langle N \cup \{k'\} \mid P' \cup \{k'\} \vdash s[k' \mapsto (t, k)] \rangle\rangle$$

if  $f @ M$  not defined,  $k' \notin N$ ,  $P' = Cl(s, P \cup \nu(v))$ .

When a non-local function is called, control is transferred to the system. In game semantics this corresponds to a *Proponent question*, and is an observable action. Following it, all the names that can be transitively reached from public names in the store also become public, so it gives both control and information to the System. Its observability is marked by a label on the transition arrow, which includes: a tag **call**, indicating that a function is called, the name of the function ( $f$ ), its arguments ( $v$ ) and a fresh continuation ( $k'$ ), which stores the code pointer; the transition also marks that part of the store which is observable because it uses publicly known names.

The counterpart rule is the system-to-program return, corresponding to a return from a non-local function.

*System-to-Program return:*

$$\langle\langle N \mid P \vdash s \rangle\rangle \xrightarrow[s']{\text{ret } v, k'} \langle N \cup \nu(v, s') \mid P \cup \nu(v, s') \vdash s[s'], f, v, k \rangle$$

if  $s(k') = (f, k)$ ,  $\nu(v, s') \cap N \subseteq P$ ,  $\lambda(\nu(v)) \subseteq \nu(s')$ ,  $s \upharpoonright \lambda(P) \sqsubseteq s'$ ,  $\kappa(s') = \emptyset$ .

This is akin to the game-semantic *Opponent answer*. Operationally it corresponds to S returning from a function. Note here that the only constraints on what S can do in this situation are *epistemic*, i.e. determined by what it *knows*:

1. it can return with any value  $v$  so long as it only contains public names or fresh names (but not *private* ones);
2. it can update any public location with any value;
3. it can return to any (public) continuation  $k'$ .

However, the part of the store which is private (i.e. with domain in  $N \setminus P$ ) cannot be modified by S. So S has no restrictions over what it can do with known names and to known names, but it cannot guess private names. Therefore it cannot do anything with or to names it does not know. The restriction on the continuation are just hygienic, as explained earlier.

There are two converse transfer rules corresponding to the program returning and the system initiating a function call:

*System-to-Program call:*

$$\langle\langle N \mid P \vdash s \rangle\rangle \xrightarrow[s']{\text{call } f, v, k} \langle N \cup \{k\} \cup \nu(v, s') \mid P \cup \{k\} \cup \nu(v, s') \vdash s[s'], f(\square), v, k \rangle$$

if  $f @ M$  defined,  $k \notin \text{dom}(s)$ ,  $\nu(f, v, s') \cap N \subseteq P$ ,  $\lambda(\nu(v)) \subseteq \nu(s')$ ,  $s \upharpoonright \lambda(P) \subseteq s'$ ,  $\kappa(s') = \emptyset$ .

*Program-to-System return:*

$$\langle N \mid P \vdash s, -, v, k \rangle \xrightarrow[s \upharpoonright \lambda(P')]{\text{ret } v, k} \langle\langle N \mid P' \vdash s \rangle\rangle, \text{ where } P' = Cl(s, P \cup \nu(v)).$$

In the case of the S-P call it is S which calls a publicly-named function from the module. As in the case of the return, the only constraint is that the function  $f$ , arguments  $v$  and the state update  $s'$  only involve public or fresh names. The hygiene conditions on the continuations impose that no continuation names are stored, for reasons already explained. Finally, the P-S return represents the action of the program yielding a final result to the system following a function call. The names used in constructing the return value are disclosed and the public part of the store is observed. In analogy with game semantics the function return is a *Proponent answer* while the system call is an *Opponent question*.

The *initial configuration* of the SLS for module  $M$  is  $S_M^0 = \langle\langle N \mid P \vdash s_0 \rangle\rangle$ . It contains a store  $s_0$  where all variables are initialised to the value specified in the declaration. The set  $N$  of names contains all the exported and imported names, all declared variables and functions. The set  $P$  contains all exported and imported names.

### 3 Compositionality

The SLS of a module  $M$  gives us an interpretation  $\llbracket M \rrbracket$  which is modular and effective (i.e. it can be executed) so no consideration of the context is required in formulating properties of modules based on their SLS. Technically, we can reason about SLS using standard tools for transition systems such as trace equivalence, bisimulation or Hennessy-Milner logic.

We first show that the SLS is consistent by proving a *compositionality* property. SLS interpretations of modules can be composed semantically in a way that is consistent with syntactic composition. Syntactic composition for modules is concatenation with renaming of un-exported function and variable names to prevent clashes, which we will denote by using  $- \cdot -$ . We call this *the principle of functional composition*.

In this section we show that we can define a semantic SLS composition  $\otimes$  so that, for an appropriate notion of bisimulation in the presence of bound names and  $\tau$ -transitions:

*Functional composition.* For any modules  $M, M'$ :  $\llbracket M \cdot M' \rrbracket \sim \llbracket M \rrbracket \otimes \llbracket M' \rrbracket$ .

Let  $\mathcal{P}$  range over program configurations, and  $\mathfrak{B}$  over system configurations. We define semantic composition of modules inductively as in Fig. 2 (all rules have symmetric, omitted counterparts). We use an extra component  $\Pi$  containing those names which have been communicated between either module and the outside system, and we use an auxiliary store  $s$  containing values of locations only. Continuation names in each  $\Pi$  are assigned Program/System polarities, thus specifying whether a continuation name was introduced by either of the modules or from the outside system. We write  $k \in \Pi_P$  when  $k \in \Pi$  has Program polarity, and dually for  $k \in \Pi_S$ . We also use the following notations for updates, where we write  $Pr$  for the set of private names  $\nu(\mathfrak{B}, \mathfrak{B}') \setminus \Pi$ .

$$\begin{aligned} (\Pi, s')^P[v, k, s] &= (\Pi', s'[s]) \text{ where } \Pi' = Cl(s'[s], \nu(v) \cup \Pi) \cup \{k\} \text{ and } k \in \Pi'_P \\ (\Pi, s')^S[v, k, s] &= (\Pi', s'[s]) \text{ where } \Pi' = \Pi \cup \nu(v, s \setminus Pr) \cup \{k\} \text{ and } k \in \Pi'_S \\ &\text{and } (s' \upharpoonright \Pi) \sqsubseteq s, s' \setminus \Pi \subseteq s, \nu(v', s \setminus Pr) \cap Pr = \emptyset \end{aligned}$$

The same notations are used when no continuation name  $k$  is included in the update. Private calls and returns are assigned  $\tau$ -labels, thus specifying the fact that they are internal transitions.

The semantic composition of modules  $M$  and  $M'$  is given by  $\llbracket M \rrbracket \otimes \llbracket M' \rrbracket = \llbracket M \rrbracket \otimes_{\Pi_0}^{s_0 \cup s'_0} \llbracket M' \rrbracket$ , where  $s_0$  is the store assigning initial values to all initial public locations of  $\llbracket M \rrbracket$ , and similarly for  $s'_0$ , and  $\Pi_0$  contains all exported and imported names. The rules of Fig. 2 feature side-conditions regarding name-privacy. These stem from nominal game semantics [10] and they guarantee that the names introduced (freshly) by  $M$  and  $M'$  do not overlap and that the names introduced by the system in the composite module do not overlap with any of the names introduced by  $M$  or  $M'$ . The former condition is necessary for correctness. The latter is typically needed for associativity.

Let us call the four participants in the composite SLS *Program A*, *System A*, *Program B*, *System B*. Whenever we use X, Y as Program or System names they can be either A or B, but different. Whenever we say *Agent* we mean Program or System. System X and Program X form *Entity X*. A state of the composite system is a pair (Agent X, Agent Y) noting that they cannot be both Programs. The composite transition rules reflect the following intuitions:



1. Internal move: 
$$\frac{\mathcal{P} \longrightarrow \mathcal{P}'}{\mathcal{P} \otimes_{\Pi}^s \mathcal{B} \longrightarrow \mathcal{P}' \otimes_{\Pi}^s \mathcal{B}} \quad \nu(\mathcal{P}') \cap \nu(\mathcal{B}) \subseteq \nu(\mathcal{P})$$
2. Cross-call: 
$$\frac{\mathcal{P} \xrightarrow[s]{\text{call } f, v, k} \mathcal{B}' \quad \mathcal{B} \xrightarrow[s]{\text{call } f, v, k} \mathcal{P}'}{\mathcal{P} \otimes_{\Pi}^{s'} \mathcal{B} \xrightarrow{\tau} \mathcal{B}' \otimes_{\Pi}^{s'[s]} \mathcal{P}'} \quad k \notin \nu(\mathcal{B})$$
3. Cross-return: 
$$\frac{\mathcal{P} \xrightarrow[s]{\text{ret } v, k} \mathcal{B}' \quad \mathcal{B} \xrightarrow[s]{\text{ret } v, k} \mathcal{P}'}{\mathcal{P} \otimes_{\Pi}^{s'} \mathcal{B} \xrightarrow{\tau} \mathcal{B}' \otimes_{\Pi}^{s'[s]} \mathcal{P}'}$$
4. Program call: 
$$\frac{\mathcal{P} \xrightarrow[s]{\text{call } f, v, k} \mathcal{B}' \quad \mathcal{B} \xrightarrow[s]{\text{call } f, v, k} \mathcal{P}'}{\mathcal{P} \otimes_{\Pi}^{s'} \mathcal{B} \xrightarrow[s'' \upharpoonright \Pi']{\text{call } f, v, k} \mathcal{B}' \otimes_{\Pi'}^{s''} \mathcal{B}} \quad (\Pi', s'') = (\Pi, s')^P[v, k, s] \quad k \notin \nu(\mathcal{B})$$
5. Program return: 
$$\frac{\mathcal{P} \xrightarrow[s]{\text{ret } v, k} \mathcal{B}' \quad \mathcal{B} \xrightarrow[s]{\text{ret } v, k} \mathcal{P}'}{\mathcal{P} \otimes_{\Pi}^{s'} \mathcal{B} \xrightarrow[s'' \upharpoonright \Pi']{\text{ret } v, k} \mathcal{B}' \otimes_{\Pi'}^{s''} \mathcal{B}} \quad (\Pi', s'') = (\Pi, s')^P[v, s]$$
6. System call: 
$$\frac{\mathcal{B} \xrightarrow[s]{\text{call } f, v, k} \mathcal{P}}{\mathcal{B} \otimes_{\Pi}^{s'} \mathcal{B}' \xrightarrow[s \upharpoonright \Pi']{\text{call } f, v, k} \mathcal{P} \otimes_{\Pi'}^{s''} \mathcal{B}'} \quad (\Pi', s'') = (\Pi, s')^S[v, k, s] \quad k \notin \nu(\mathcal{B}') \setminus \Pi_S$$
7. System return: 
$$\frac{\mathcal{B} \xrightarrow[s]{\text{ret } v, k} \mathcal{P}}{\mathcal{B} \otimes_{\Pi}^{s'} \mathcal{B}' \xrightarrow[s \upharpoonright \Pi']{\text{ret } v, k} \mathcal{P} \otimes_{\Pi'}^{s''} \mathcal{B}'} \quad (\Pi', s'') = (\Pi, s')^S[v, s] \quad k \notin \nu(\mathcal{B}')$$

**Fig. 2.** Rules for semantic composition

- Rule 1: If Program X makes an internal (operational) transition System Y is not affected.
- Rules 2-3: If Program X makes a system transition to System X and System Y can match the transition going to Program Y then the composite system makes an internal transition. This is the most important rule and it is akin to game semantic composition via “synchronisation and hiding”. It signifies  $M$  making a call (or return) to (from) a function present in  $M'$ .
- Rules 4-5: If Program X makes a system transition that cannot be matched by Entity Y then it is a system transition in the composite system, a non-local call or return.
- Rules 6-7: From a composite system configuration (both entities are in a system configuration) either Program X or Program Y can become active via a call or return from the system.

**Lemma 1.** *Let  $X_1 \otimes_{\Pi}^s X_2$  be a state in the transition graph of  $\llbracket M \rrbracket \otimes \llbracket M' \rrbracket$  that is reachable from the initial state. Then, if each  $X_i$  includes the triple  $(N_i, P_i, s_i)$ , the following conditions hold.*

- $(N_1 \setminus P_1) \cap N_2 = N_1 \cap (N_2 \setminus P_2) = \emptyset$ ,  $P_1 \setminus \Pi = P_2 \setminus \Pi$ ,  $\Pi \subseteq \nu(s) \cup \kappa(P_1 \cup P_2) \subseteq P_1 \cup P_2$  and  $\nu(\text{dom}(s)) = \lambda(P_1 \cup P_2)$ .
- $\text{dom}(\kappa(s_1)) \cap \text{dom}(\kappa(s_2)) = \emptyset$ ,  $(\text{dom}(\kappa(s_1)) \cup \text{dom}(\kappa(s_2))) \cap \Pi_S = \emptyset$  and  $\kappa(P_1 \cap P_2) \setminus \Pi_S = \kappa(P_1 \cup P_2) \setminus \Pi$ .
- If both  $X_1, X_2$  are system configurations and  $X_1 \otimes_{\Pi}^s X_2$  is preceded by a state of the form  $\mathcal{P} \otimes_{\Pi'}^{s'} \mathcal{B}$  then  $s \upharpoonright P_1 \subseteq s_1$  and  $s \upharpoonright (P_2 \setminus P_1) \subseteq s_2$ , and dually if preceded by  $\mathcal{B} \otimes_{\Pi'}^{s'} \mathcal{P}$ . Thus, in both cases,  $s \upharpoonright (P_i \setminus (P_1 \cap P_2)) \subseteq s_i$  for  $i = 1, 2$ .
- Not both  $X_1, X_2$  are program configurations. If  $X_i$  is a program configuration then  $s \upharpoonright (P_{3-i} \setminus P_i) \subseteq s_{3-i}$ .

Semantic composition introduces a notion of private names: internal continuation names passed around between the two modules in order to synchronise their mutual function calls. As the previous lemma shows, these names remain private throughout the computation. Therefore, in checking bisimilarity for such reduction systems, special care has to be taken for these private names so that external system transitions capturing them do not affect these checks. This is standard procedure in calculi with name-binding.

We define the following translation  $R$  from reachable composite states of  $\llbracket M \rrbracket \otimes \llbracket M' \rrbracket$  to states of  $\llbracket M \cdot M' \rrbracket$ .

$$\begin{aligned} \langle\langle N_1 \mid P_1 \vdash s_1 \rangle\rangle \otimes_{\Pi}^s \langle\langle N_2 \mid P_2 \vdash s_2 \rangle\rangle &\longmapsto \langle\langle (N_1 \cup N_2) \setminus K \mid \Pi \vdash (\hat{s}_1[s'] \cup \hat{s}_2[s']) \setminus K \rangle\rangle \\ \langle\langle N_1 \mid P_1 \vdash s_1 \rangle\rangle \otimes_{\Pi}^s \langle N_2 \mid P_2 \vdash s_2, t, v, k \rangle &\longmapsto \langle\langle (N_1 \cup N_2) \setminus K \mid \Pi \vdash \hat{s}_1[\hat{s}_2] \setminus K, t', v, k' \rangle\rangle \\ \langle N_1 \mid P_1 \vdash s_1, t, v, k \rangle \otimes_{\Pi}^s \langle\langle N_2 \mid P_2 \vdash s_2 \rangle\rangle &\longmapsto \langle\langle (N_1 \cup N_2) \setminus K \mid \Pi \vdash \hat{s}_2[\hat{s}_1] \setminus K, t', v, k' \rangle\rangle \end{aligned}$$

where  $K = \kappa(P_1 \cap P_2) \setminus \Pi_S$ ,  $s' = s \upharpoonright (P_1 \cap P_2)$ ,  $\hat{s}_i = s_i[k \mapsto (s_1, s_2)_K(s_i(n))]$  for all  $k \in \text{dom}(\kappa(s_i))$ , and  $(t', k') = (s_1, s_2)_K(t, k)$ . The function  $(s_1, s_2)_K$  fetches the full external frame stack and the external continuation searching back from  $(t, k)$ , that is,  $(s_1, s_2)_K(t, k) = (t, k)$  if  $k \notin K$ , otherwise if  $k \in K$  and  $s_i(k) = (t', k')$  then  $(s_1, s_2)_K(t' \circ t, k')$ .

The translation merges names from the component configurations and deletes the names in  $K$ : these private names do not appear in  $\llbracket M \cdot M' \rrbracket$ , as there the corresponding function calls happen without using the call-return mechanism. It also sets  $\Pi$  as the set of public names. Moreover, the total store is computed as follows. In system configurations we just take the union of the component stores and update them with the values of  $s$ , which contains the current values of all common public names. In program configurations we use the fact that the P-component contains more recent values than those of the S-component.

**Proposition 2.** *For  $R$  defined as above and  $X_1 \otimes_{\Pi}^s X_2$  a reachable configuration,*

1. if  $X_1 \otimes_{\Pi}^s X_2 \xrightarrow{\tau} X'_1 \otimes_{\Pi}^{s'} X'_2$  then  $R(X_1 \otimes_{\Pi}^s X_2) = R(X'_1 \otimes_{\Pi}^{s'} X'_2)$ ,
2. if  $X_1 \otimes_{\Pi}^s X_2 \rightarrow X'_1 \otimes_{\Pi}^s X'_2$  then  $R(X_1 \otimes_{\Pi}^s X_2) \rightarrow R(X'_1 \otimes_{\Pi}^s X'_2)$ ,
3. if  $R(X_1 \otimes_{\Pi}^s X_2) \rightarrow Y$  and  $X_1 \otimes_{\Pi}^s X_2 \not\rightarrow$  then  $X_1 \otimes_{\Pi}^s X_2 \rightarrow X'_1 \otimes_{\Pi}^s X'_2$  with  $Y = R(X'_1 \otimes_{\Pi}^s X'_2)$ ,
4. if  $X_1 \otimes_{\Pi}^s X_2 \xrightarrow[\alpha]{s'} X'_1 \otimes_{\Pi}^{s''} X'_2$  then  $R(X_1 \otimes_{\Pi}^s X_2) \xrightarrow[\alpha]{s'} R(X'_1 \otimes_{\Pi}^{s''} X'_2)$ ,
5. if  $R(X_1 \otimes_{\Pi}^s X_2) \xrightarrow[\alpha]{s'} Y$ ,  $X_1 \otimes_{\Pi}^s X_2 \not\rightarrow$  and  $\nu(\alpha) \cap K = \emptyset$  then  $X_1 \otimes_{\Pi}^s X_2 \xrightarrow[\alpha]{s'} X'_1 \otimes_{\Pi}^{s''} X'_2$  with  $Y = R(X'_1 \otimes_{\Pi}^{s''} X'_2)$ ,

where  $K$  is obtained from  $X_1 \otimes_{\Pi}^s X_2$  as above.

With bisimilarity between semantic and syntactic composites defined as above, functional composition follows immediately as a consequence.

## 4 Reasoning about SLS

The epistemically-constrained system-level semantics gives a security-flavoured semantics for the programming language which is reflected by its logical properties and by the notion of equivalence it gives rise to.

We will see that certain properties of traces in the SLS of a module correspond to “secrecy violations”, i.e. undesirable disclosures of names that are meant to stay secret. In such traces it is reasonable to refer to the System as an *attacker* and consider its actions an attack. We will see that although the attack cannot be realised within the given language it can be enacted in a realistic system by system-level actions.

We will also see that certain equivalences that are known to hold in conventional semantics still hold in a system-level model. This means that even in the presence of an omnipotent attacker, unconstrained by a prescribed set of language constructs, the epistemic restrictions can prevent certain observations, not only by the programming context but by any ambient computational system. This is a very powerful notion of equivalence which embodies *tamper-resistance* for a module.

*Note.* We chose these examples to illustrate the conceptual interest of the SLS-induced properties rather than as an illustration of the mathematical power of SLS-based reasoning techniques. For this reason, we chose examples which are as simple and clear as possible.

### 4.1 A system-level attack: violating secrecy

This example is inspired by a flawed security protocol which is informally described as follows.

Consider a secret, a locally generated key and an item of data read from the environment. If the local key and the input data are equal then output the secret, otherwise output the local key.

In a conventional process-calculus syntax the protocol can be written as

$$\nu s \nu k. \text{in}(a). \text{if } k=a \text{ then out}(s) \text{ else out}(k).$$

It is true that the secret  $s$  is not leaked because the local  $k$  cannot be known as it is disclosed only at the very end. This can be proved using bisimulation-based techniques for anonymity. Let us consider an implementation of the protocol:

$$\begin{aligned}
& \langle N_0 \mid P_0 \vdash \emptyset \rangle \xrightarrow[\emptyset]{\text{call } \text{prot}() , k} \langle N_0, k \mid P_0, k \vdash \emptyset, -, E, k \rangle \\
& \longrightarrow^* \langle N_1, k, a_0, a_1 \mid P_0, k \vdash (\mathbf{s} \mapsto a_0, p \mapsto a_1, \mathbf{x} \mapsto 0), \\
& \quad (\square; \text{if}(*x == *p) \text{ then } *s \text{ else } *p) \circ (\mathbf{x} = \square) \circ (\text{read}(\square)), (), k \rangle \\
& \xrightarrow[\emptyset]{\text{call } \text{read}() , k'} \langle \langle N_1, k, k', a_0, a_1 \mid P_1 \vdash (\mathbf{s} \mapsto a_0, \mathbf{k} \mapsto a_1, \mathbf{x} \mapsto 0, k' \mapsto (t, k)) \rangle \rangle \\
& \xrightarrow[\emptyset]{\text{ret } a_2, k'} \langle N_2 \mid P_1, a_2 \vdash (\mathbf{s} \mapsto a_0, \mathbf{k} \mapsto a_1, \mathbf{x} \mapsto 0, k' \mapsto (t, k)), t, a_2, k \rangle \\
& \longrightarrow^* \langle N_2 \mid P_1, a_2 \vdash (\mathbf{s} \mapsto a_0, \mathbf{k} \mapsto a_1, \mathbf{x} \mapsto a_2, k' \mapsto (t, k)), -, a_1, k \rangle \\
& \xrightarrow[\emptyset]{\text{ret } a_1, k} \langle \langle N_2 \mid P_2, a_2, a_1 \vdash (\mathbf{s} \mapsto a_0, \mathbf{k} \mapsto a_1, \mathbf{x} \mapsto a_2, k' \mapsto (t, k)) \rangle \rangle \\
& \xrightarrow[\emptyset]{\text{ret } a_1, k'} \langle N_2 \mid P_1, a_2, a_1 \vdash (\mathbf{s} \mapsto a_0, \mathbf{k} \mapsto a_1, \mathbf{x} \mapsto a_2, k' \mapsto (t, k)), t, a_1, k \rangle \\
& \longrightarrow^* \langle N_2 \mid P_1, a_2, a_1 \vdash (\mathbf{s} \mapsto a_0, \mathbf{k} \mapsto a_1, \mathbf{x} \mapsto a_1, k' \mapsto (t, k)), -, a_0, k \rangle \\
& \xrightarrow[\emptyset]{\text{ret } a_0, k} \langle \langle N_2 \mid P_2, a_2, a_1, a_0 \vdash (\mathbf{s} \mapsto a_0, \mathbf{k} \mapsto a_1, \mathbf{x} \mapsto a_2, k' \mapsto (t, k)) \rangle \rangle.
\end{aligned}$$

Above,  $t = (\square; \text{if}(*x == *k) \text{ then } *s \text{ else } *k) \circ (\mathbf{x} = \square)$ ,  $N_0 = P_0 = \{\text{prot}, \text{read}\}$ ,  $N_1 = N_0 \cup \{\mathbf{s}, \mathbf{k}, \mathbf{x}\}$ ,  $N_2 = N_1 \cup \{k, k', a_0, a_1, a_2\}$  and  $P_1 = P_0 \cup \{k, k'\}$ .

**Fig. 3.** Secret  $a_0$  leaks.

```

export prot;
import read;
decl prot( ) { local s, k, x;
               s = new(); k = new(); x = read();
               if (*x == *k) then *s else *k }

```

We have local variables  $\mathbf{s}$  holding the “secret location” and  $\mathbf{k}$  holding the “private location”. We use the non-local, system-provided, function `read` to obtain a name from the system, which cannot be that stored at  $\mathbf{s}$  or  $\mathbf{k}$ . A value is read into  $\mathbf{x}$  using untrusted system call `read()`. Can the secrecy of  $\mathbf{s}$  be violated by making the name stored into it public? Unlike in the process-calculus model, the answer is “yes”.

The initial state is  $\langle \langle \text{prot}, \text{read} \mid \text{prot}, \text{read} \vdash \emptyset \rangle \rangle$ . We denote the body of `prot` by  $E$ . The transition corresponding to the secret being leaked is shown in Fig. 3. The labelled transitions are the interactions between the program and the system and are interpreted as follows:

1. system calls `prot()` giving continuation  $k$
2. program calls `read()` giving fresh continuation  $k'$
3. system returns (from `read`) using  $k'$  and producing fresh name  $a_2$
4. program returns (from `prot`) leaking local name  $a_1$  stored in  $\mathbf{k}$
5. system uses  $k'$  to fake a second return from `read`, using the just-learned name  $a_1$  as a return value

6. with  $a_1$  the program now returns the secret  $a_0$  stored in  $s$  to the environment. Values of  $a_2$  are omitted as they do not affect the transitions.

The critical step is (5), where the system is using a continuation in a presumably illegal, or at least unexpected, way. This attack could be executed in a language with `call-cc`-like control features, but these are lacking from our language. We do not even need a richer ambient language to show how a *system-level* attack can be actually implemented. Surprisingly, all we need is an implementation of `read()` which will wait to receive a value from the attacker, and a `main()` function which calls `prot()` and reports the value.

We execute the attack by running the (closed) program in a *virtual machine* in the following way:

1. execute the program normally until the `read` function is called;
2. pause the virtual machine, save its state and exit;
3. duplicate the file storing the state of the virtual machine and re-start one instance of the virtual machine;
4. feed an arbitrary value to `read()`;
5. when the program terminates normally remember the final value, which corresponds to  $a_1$ , stored in  $k$ ;
6. re-start the other instance of the virtual machine;
7. feed  $a_1$  to `read()`;
8. when this instance of the program terminates normally it leaks the secret  $a_0$  from  $s$ .

Note that the interaction between the attacker and the two instances of the program correspond precisely to the labelled actions in the attack.

What is remarkable about this attack is that both the term and the context are written in a simple programming language that cannot implement the attack! The attack happens because of a system-level action, the cloning of a virtual machine. Also note that this is not a theoretical attack. Our language is a subset of C and it can be compiled, with small syntactic adjustments, by conventional C compilers and executed on conventional operating systems. Any virtualisation platform such as VMWare or VirtualBox can be used to express this attack.

## 4.2 Equivalence

Functional Compositionality gives an internal consistency check for the semantics. This already shows that our language is “well behaved” from a system-level point of view. In this section we want to further emphasise this point. We can do that by proving that there are nontrivial equivalences which hold. There are many such equivalences we can show, but we will choose a simple but important one, because it embodies a principle of locality for state.

This deceptively simple example was first given in [12] and establishes the fact that a *local* variable cannot be interfered with by a *non-local* function. This was an interesting example because it highlighted a significant shortcoming of *global state* models of imperative programming. Although not pointed out at the

time, functor-category models of state developed roughly at the same time gave a mathematically clean solution for this equivalence, which followed directly from the type structure of the programming language [16].

For comparing SLS LTSs we can use a simpler notion of bisimulation which relates configurations, and modules, that have common public names.

**Definition 3.**  $\mathcal{R}$  is a simulation if, whenever  $(X_1, X_2) \in \mathcal{R}$ ,

- $X_1$  and  $X_2$  have the same public names;
- $X_1 \rightarrow X'_1$  implies  $(X'_1, X_2) \in \mathcal{R}$ ;
- $X_1 \xrightarrow{s} X'_1$  implies  $(\pi \cdot X_2) \xrightarrow{s} X'_2$  and  $(X'_1, X'_2) \in \mathcal{R}$ , for some name permutation  $\pi$  such that  $\pi(a) = a$  for all public names  $a$  of  $X_1$  and  $X_2$ .

$\mathcal{R}$  is a bisimulation if it and its inverse are simulations. We say that modules  $M_1$  and  $M_2$  are bisimilar if there is a bisimulation  $\mathcal{R}$  such that  $(S_{M_1}^0, S_{M_2}^0) \in \mathcal{R}$ .

**Proposition 4.** Bisimulation is a congruence for module composition  $- \cdot -$ .

The proof uses the reduction of syntactic composition to semantic composition then uses Prop. 2 to show that bisimulation is preserved by semantic composition with the same module, which is immediate. This is unsurprising, since system-level bisimilarity is more fine-grained than contextual equivalence in the programming language.

It is straightforward to check that the following three programs have bisimilar SLS transition systems:

```
export f; import g; decl f() {local x; g(); return *x;}
export f; import g; decl x; decl f() {g(); return *x;}
export f; import g; decl f() {g(); return 0;}
```

Intuitively, the reason is that in the first two programs *f*-local (module-local, respectively) variable *x* is never visible to non-local function *g*, and will keep its initial value, which is 0. The bisimulation relation is straightforward as the three LTSs are equal modulo silent transitions and permutation of private names for *x*. Other equivalences, for example in the style of *parametricity* [13] also hold, with simple proofs of equivalence via bisimulation.

## 5 Conclusion

The Dolev-Yao-like characterisation of the Opponent in this semantics suggests that this is a model suitable for modelling security properties. The system-level semantics presupposes certain strong guarantees of secrecy and integrity for the combined execution environment consisting of compiler and operating system: certain location names must be kept secret; the Program source code cannot be altered; even if the name of a function or continuation is disclosed the names used in the function and in the continuation remain secret. All these requirements can be gathered under the principle that *the System can make no low-level attacks against the Program*. This is also consistent with the Dolev-Yao principle that the attacker can manipulate messages, but without breaking cryptography.

Compilers such as `gcc` do not implement a system-level semantics since locations are not secret and the code layout is known, allowing low-level attacks. Most security violation of C code are through low-level attacks such as buffer overflows. However, there are significant research and industrial efforts to produce tamper-proof code through techniques such as address layout randomisation [7,14], address obfuscation [3], instruction-set obfuscation [8] or secure processors [15]. A system-level semantics gives a semantically-directed specification for a tamper-proof compiler. System-level semantics also gives a basis for the study of security properties of programs compiled with such tamper-proof compilers, highlighting *logical* attacks, such as the secrecy violation in Sec. 4.1. The system does not *guess* any of the secrets of the program and it does not tamper with its code, but it clones it wholesale then plays the two instances against each other, a typical replay attack. Conversely, the equivalences of Sec. 4.2 present opportunities for optimisations which hold not only relative to certain compilers, but to a more comprehensive concept of execution environment.

## References

1. S. Abramsky, R. Jagadeesan, and P. Malacaria. Full abstraction for PCF. *Inf. Comput.*, 163(2):409–470, 2000.
2. S. Abramsky and G. McCusker. Linearity, sharing and state: a fully abstract game semantics for Idealized Algol with active expressions. *ENTCS*, 3, 1996.
3. S. Bhatkar, D. C. DuVarney, and R. Sekar. Address obfuscation: an efficient approach to combat a board range of memory error exploits. *USENIX*, 2003.
4. M. Gabbay and A. M. Pitts. A new approach to abstract syntax involving binders. In *LICS*, pages 214–224, 1999.
5. D. R. Ghica and A. Murawski. Angelic semantics of fine-grained concurrency. *Annals of Pure and Applied Logic*, 151(2-3):89–114, 2008.
6. J. M. E. Hyland and C.-H. L. Ong. On full abstraction for PCF: I, II, and III. *Inf. Comput.*, 163(2):285–408, 2000.
7. R. Jagadeesan, C. Pitcher, J. Rathke, and J. Riely. Local memory via layout randomization. In *CSF*, pages 161–174, 2011.
8. G. S. Kc, A. D. Keromytis, and V. Prevelakis. Countering code-injection attacks with instruction-set randomization. In *CCS*, pages 272–280, 2003.
9. J. Laird. Full abstraction for functional languages with control. In *LICS*, 1997.
10. J. Laird. A game semantics of local names and good variables. In *FoSSaCS*, 2004.
11. J. Laird. A fully abstract trace semantics for general references. In *ICALP*, 2007.
12. A. R. Meyer and K. Sieber. Towards fully abstract semantics for local variables. In *POPL*, pages 191–203, 1988.
13. P. W. O’Hearn and R. D. Tennent. Parametricity and local variables. *J. ACM*, 42(3):658–709, 1995.
14. H. Shacham, M. Page, B. Pfaff, E.-J. Goh, N. Modadugu, and D. Boneh. On the effectiveness of address-space randomization. In *CCS*, pages 298–307, 2004.
15. G. E. Suh, D. Clarke, B. Gassend, M. v. Dijk, and S. Devadas. Efficient memory integrity verification and encryption for secure processors. In *MICRO*, 2003.
16. R. D. Tennent. Semantical analysis of specification logic. *Inf. Comput.*, 85(2):135–162, 1990.

## A Nominal Sets

It is handy to introduce here some basic notions from the theory of nominal sets [4]. We call *nominal structure* any structure which may contain names, i.e. elements of  $\mathcal{N}$ , and we denote by  $Perm$  the set of finite permutations on  $\mathcal{N}$  which are sort-preserving (i.e. if  $a \in \mathcal{N}_\lambda$  then  $\pi(a) \in \mathcal{N}_\lambda$ , etc.). For example,  $\text{id} = \{(a, a) \mid a \in \mathcal{N}\} \in Perm$ . For each set  $X$  of nominal structures of interest, we define a function  $\pi \cdot \_ : Perm \times X \rightarrow X$  such that  $\pi \cdot (\pi' \cdot x) = (\pi \circ \pi') \cdot x$  and  $\text{id} \cdot x = x$ , for all  $x \in X$  and  $\pi, \pi' \in Perm$ .  $X$  is called a *nominal set* if all its elements involve finitely many names, that is, for all  $x \in X$  there is a finite set  $S \subseteq \mathcal{N}$  such that  $\pi \cdot x = x$  whenever  $\forall a \in S. \pi(a) = a$ . The minimal such set  $S$  is called the *support* of  $x$  and denoted by  $\nu(x)$ . For example,  $\mathcal{N}$  is a nominal set with action  $\pi \cdot a = \pi(a)$ , and so is  $\mathcal{P}_{\text{fn}}(\mathcal{N})$  with action  $\pi \cdot S = \{\pi(a) \mid a \in S\}$ .

Also, any set of non-nominal structures is a nominal set with trivial action  $\pi \cdot x = x$ . More interestingly, if  $X, Y$  are nominal sets then so is  $X \times Y$  with action  $\pi \cdot (x, y) = (\pi \cdot x, \pi \cdot y)$ . This extends to arbitrary products and to strings. Also, if  $X$  is a nominal set then so is the set  $\bigcup_{n \in \omega} (\{1, \dots, n\} \rightarrow X)$  with action  $\pi \cdot f = \{(i, \pi \cdot x) \mid (i, x) \in f\}$ . Finally, if  $X, Y$  are nominal sets then so is the set  $X \rightarrow_{\text{fn}} Y$  with action  $\pi \cdot f = \{(\pi \cdot x, \pi \cdot y) \mid (x, y) \in f\}$ .

## B Proof of Proposition 2

*Proof.* For 1, let  $X_1 = \langle N_1 \mid P_1 \vdash s_1, t \circ f(\square), v, k \rangle$ ,  $X_2 = \langle N_2 \mid P_2 \vdash s_2 \rangle$  and the  $\tau$ -transition being due to an internal transition with label  $(s_i, \text{call } f, v, k')$ . Thus,  $X'_1 = \langle N'_1 \mid P'_1 \vdash s'_1 \rangle$ ,  $X'_2 = \langle N'_2 \mid P'_2 \vdash s'_2, f(\square), v, k' \rangle$ , and so  $R(X_1 \otimes_{\Pi}^s X_2) = \langle N_0 \mid \Pi \vdash s_0, t_0, v, k_0 \rangle$  and  $R(X'_1 \otimes_{\Pi}^s X'_2) = \langle N'_0 \mid \Pi \vdash s'_0, t'_0, v, k'_0 \rangle$ . Computing  $K, K'$  as above, we have  $K' = K \cup \{k'\}$ . Moreover,  $s'_1 = s_1[k' \mapsto (t, k)]$  and  $s'_2 = s \cup (s_2 \setminus P_2)$ , so  $(t'_0, k'_0) = (s'_1, s'_2)_{K'}(f(\square), k') = (s'_1, s'_2)_{K'}(t \circ f(\square), k)$ . Since  $k'$  is fresh,  $(s'_1, s'_2)_{K'}(t \circ f(\square), k) = (s_1, s_2)_K(t \circ f(\square), k) = (t_0, k_0)$ . Moreover,  $N_0 = (N_1 \cup N_2) \setminus K$  and  $N'_0 = (N'_1 \cup N'_2) \setminus K' = (N_1 \cup \{k'\} \cup N_2 \cup \nu(v, s_i)) \setminus K'$ . As  $\nu(v, s_i) \subseteq N_1$  and  $k' \in K'$ , we get  $N_0 = N'_0$ . Finally,  $s_0 = \hat{s}_2[\hat{s}_1] \setminus K$  and  $s'_0 = \hat{s}'_1[\hat{s}'_2] \setminus K'$ . Thus,  $s'_0 = \hat{s}_1[\hat{s}'_2] \setminus K = \hat{s}_1[s' \cup (\hat{s}_2 \setminus \lambda(P_2))] \setminus K$ . Moreover,  $s' = s_1 \upharpoonright \lambda(P'_1)$  so  $s'_0 = \hat{s}_1[\hat{s}_2 \setminus \lambda(P_2)] \setminus K$ . But now note that  $\text{dom}(s_2 \setminus \lambda(P_2)) \cap \text{dom}(s_1) = \emptyset$ : by the previous lemma,  $\text{dom}(s_1)$  and  $\text{dom}(s_2)$  share no continuation names, and if  $a$  is a location name in  $\text{dom}(s_2) \setminus P_2$  then  $a \notin N_1$ . Thus,  $s_0 = s'_0$ . Similarly if the  $\tau$ -transition is due to an internal return.

Item 2 is straightforward. For 3, the only interesting issue is establishing that if  $X_1 \otimes_{\Pi}^s X_2$  is in such a form that a  $\tau$ -transition needs to take place then the latter is possible. This follows directly from the definition of the transitions and the conditions of the previous lemma. In the following cases we consider call transitions; cases with return transitions are treated in a similar manner.

For 4, let  $X_1 = \langle N_1 \mid P_1 \vdash s_1 \rangle$ ,  $X_2 = \langle N_2 \mid P_2 \vdash s_2 \rangle$ ,  $\alpha = (s', \text{call } f, v, k)$  and suppose the transition is due to  $X_1$  reducing to  $X'_1 = \langle N'_1 \mid P'_1 \vdash s'_1, f(\square), v, k \rangle$  with label  $(s_i, \text{call } f, v, k)$ . We have  $\Pi' = \Pi \cup \nu(v, k, s_i \setminus Pr)$ ,  $Pr = (N_1 \cup N_2) \setminus \Pi$ ,  $s' = s_i \upharpoonright \Pi'$  and  $\nu(v, s_i \setminus Pr) \cap Pr = \emptyset$ . Let  $R(X_1 \otimes_{\Pi}^s X_2) = \langle N_0 \mid \Pi \vdash s_0 \rangle$ . As



$k \notin \text{dom}(s_1)$  and  $k \notin \nu(X_2) \setminus \Pi_S$ , by previous lemma we obtain  $k \notin \text{dom}(s_0)$ , so the latter reduces to  $\langle N'_0 \mid P \vdash s'_0, f(\square), v, k \rangle$  with transition  $(s''', \text{call } f, v, k)$ , for any appropriate  $s'''$ . In fact, if  $\nu(v, s') \cap N_0 \subseteq \Pi$  then we can choose  $s''' = s'$ . Indeed,  $(\nu(v, s') \cap N_0) \setminus \Pi \subseteq \nu(v, s') \cap (N_0 \setminus \Pi) \subseteq \nu(v, s') \cap Pr = \nu(v, s_i \upharpoonright \Pi') \cap Pr = \nu(v, s_i \setminus Pr) \cap Pr = \emptyset$ . Let  $R(X'_1 \otimes_{\Pi'}^s X_2) = \langle N''_0 \mid \Pi' \vdash s''_0 \rangle$ . We can see that  $N'_0 = N''_0$ . Also,  $P = \Pi \cup \{k\} \cup \nu(v, s')$  while  $\Pi' = \Pi \cup \nu(v, k, s_i \setminus Pr) = \Pi \cup \nu(v, k, s')$ . Moreover,  $s'_0 = s' \cup (s_0 \setminus \lambda(\Pi)) = s' \cup ((\hat{s}_1[s_{12}] \cup \hat{s}_2[s_{12}]) \setminus (K \cup \lambda(\Pi)))$  with  $s_{12} = s \upharpoonright (P_1 \cap P_2)$ , and  $s''_0 = \hat{s}_2[\hat{s}'_1] \setminus K' = \hat{s}_2[s_i \cup (\hat{s}_1 \setminus \lambda(P_1))] \setminus K'$ . Note that  $K' = K$ . Moreover,  $s'_0$  and  $s''_0$  agree on the domain of  $s'$  and on continuation names. Also, if location name  $a \in N'_0 \setminus N_0$  then  $a \in \nu(v, s')$  and thus  $a \in \text{dom}(s')$ . Thus, we need to show that  $s'_0, s''_0$  agree on location names  $a$  from  $N_0 \setminus \Pi$ . If  $a \in N_1 \setminus P_1$  then  $s'_0(a) = s_1(a) = s''_0(a)$ , and similarly if in  $N_2 \setminus P_2$  using the fact that  $(N_2 \setminus P_2) \cap N_1 = \emptyset$ . Finally, if  $a \in P_1 \setminus \Pi = P_2 \setminus \Pi$  then  $s'_0(a) = s(a) = s_i(a) = s''_0(a)$ , by restrictions on  $s_i$ .

Now let  $X_1 = \langle N_1 \mid P_1 \vdash s_1, t \circ f(\square), v, k \rangle$ ,  $X_2 = \langle N_2 \mid P_2 \vdash s_2 \rangle$ ,  $\alpha = \text{call } f, v, k'$  and suppose the transition is due to  $X_1$  reducing to  $X'_1 = \langle N'_1 \mid P'_1 \vdash s'_1 \rangle$  with label  $(s_i, \text{call } f, v, k')$ . We have  $(\Pi', s'') = (\Pi, s)[v, s_i]$  and  $s' = s'' \upharpoonright \Pi'$ . We can assume, by definition, that  $(s_1, s_2)_K(t \circ f(\square), k) = (t_0 \circ f(\square), k_0)$ , so  $R(X_1 \otimes_{\Pi}^s X_2) = \langle N_0 \mid \Pi \vdash s_0, t_0 \circ f(\square), v, k_0 \rangle$ . As  $f$  is not defined in either of the modules and  $k'$  is completely fresh, the latter reduces to  $\langle N'_0 \mid P \vdash s'_0 \rangle$  with transition  $(s''', \text{call } f, v, k')$ . Let  $R(X'_1 \otimes_{\Pi'}^s X_2) = \langle N''_0 \mid \Pi' \vdash s''_0 \rangle$ . It is easy to see that  $N'_0 = N''_0$ . Moreover,  $s'_0 = s_0[k' \mapsto (t_0, k_0)]$  and  $s''_0 = (\hat{s}'_1[s''_{12}] \cup \hat{s}_2[s''_{12}]) \setminus K$  where  $s''_{12} = s'' \upharpoonright (P'_1 \cap P_2)$ . Note that  $K' = \kappa(P'_1 \cap P_2) = K$  and  $s''_0(k') = \hat{s}'_1(k') = (s'_1, s_2)_{K'}(t, k) = (t_0, k_0)$ . Also,  $s'_0$  and  $s''_0$  agree on all other continuation names. Thus, in order to establish that  $s'_0 = s''_0$ , it suffices to show that  $s_2[s_1]$  and  $s_1[s''_{12}] \cup s_2[s''_{12}]$  agree on locations. From the previous lemma,  $s''$  agrees with  $s_1$  on locations in  $P'_1$  and with  $s_2$  on locations in  $P_2 \setminus P'_1$ , and so  $s'' \subseteq s_2[s_1]$ . Thus,  $\lambda(s_1[s''_{12}] \cup s_2[s''_{12}]) = \lambda(s_1 \cup (s_2 \setminus P'_1)) = \lambda(s_2[s_1])$ .

For public names, we have  $P = Cl(s_0, \Pi \cup \nu(v)) \cup \{k'\} = Cl(s'_0, \Pi \cup \nu(v, k'))$  while  $\Pi' = Cl(s'', \Pi \cup \nu(v, k'))$ . As  $\kappa(P) = \kappa(\Pi) \cup \{k'\} = \kappa(\Pi')$ , we can focus on location names. We have  $s'' \subseteq s'_0$  and, moreover,  $\text{dom}(s'') = \lambda(P'_1 \cup P_2) \supseteq \lambda(\Pi \cup \nu(v, k'))$ , thus  $P = \Pi'$ . Finally,  $s' = s'''$  follows from the fact that these are restrictions of the final stores to the final sets of public location names.

For 5, let  $X_1 = \langle N_1 \mid P_1 \vdash s_1 \rangle$ ,  $X_2 = \langle N_2 \mid P_2 \vdash s_2 \rangle$ ,  $R(X_1 \otimes_{\Pi}^s X_2) = \langle N_0 \mid \Pi \vdash s_0 \rangle$  and  $\alpha = (s', \text{call } f, v, k)$ . We have that  $f$  is defined in  $M \cdot M'$  so WLOG assume that it is defined in  $M$ . Then,  $X_1$  reduces to  $X'_1 = \langle N'_1 \mid P'_1 \vdash s'_1, f(\square), v, k \rangle$  with  $(s_i, \text{call } f, v, k)$ ,  $s_i = s' \cup (s \setminus \Pi)$ , if the relevant conditions for S-P calls are satisfied.

If  $k \in \text{dom}(s_1)$  then, by lemma,  $k \notin \Pi_S$ . By assumption,  $k \in \Pi$  so  $k \notin \kappa(P_1 \cup P_2) \setminus \Pi$  and thus, by lemma,  $k \notin \kappa(P_1 \cap P_2) \setminus \Pi_S$  so  $k \notin P_2$ . But the latter would imply  $k \in \text{dom}(s_0)$ , which is disallowed by definition. Thus,  $k \notin \text{dom}(s_1)$ .

Moreover, if  $a \in \nu(v, s_i) \cap (N_1 \setminus P_1) = \nu(v, s') \cap (N_1 \setminus P_1)$  then  $a \in \nu(v, s') \cap (N_0 \setminus P_1)$  and  $a \notin P_2$ , so  $a \in \nu(v, s') \cap (N_0 \setminus (P_1 \cup P_2)) \subseteq \nu(v, s') \cap (N_0 \setminus \Pi)$ , thus contradicting the conditions for the transition  $\alpha$ . We still need to check that  $s_1 \upharpoonright \lambda(P_1) \subseteq s_i = s' \cup (s \setminus \Pi)$ . Given that  $s_0 \upharpoonright \lambda(\Pi) = (s_1[s_{12}] \cup s_2[s_{12}]) \upharpoonright$

$\lambda(\Pi) \sqsubseteq s'$ , the condition follows from the previous lemma. We therefore obtain a transition from  $X_1 \otimes_{\Pi}^s X_2$  to  $X'_1 \otimes_{\Pi'}^{s''} X_2$ ; the relevant side-conditions are shown to be satisfied similarly as above. Finally, working as in 4, we obtain  $R(X'_1 \otimes_{\Pi'}^{s''} X_2) = Y$  and  $s' = s'''$ .

Now let  $X_1 = \langle N_1 \mid P_1 \vdash s_1, t \circ f(\square), v, k \rangle$ ,  $X_2 = \langle \langle N_2 \mid P_2 \vdash s_2 \rangle \rangle$ ,  $R(X_1 \otimes_{\Pi}^s X_2) = \langle N_0 \mid \Pi \vdash s_0, t_0 \circ f(\square), v, k_0 \rangle$  and  $\alpha = \mathbf{call} f, v, k'$ . By hypothesis,  $k'$  is fresh and therefore  $X_1$  reduces to  $X'_1 = \langle \langle N'_1 \mid P'_1 \vdash s'_1 \rangle \rangle$  with  $(s_i, \mathbf{call} f, v, k')$ , and thus  $X_1 \otimes_{\Pi}^s X_2$  reduces to  $X'_1 \otimes_{\Pi'}^{s''} X_2$  with  $(s''', \mathbf{call} f, v, k')$ . Working as in 4,  $R(X'_1 \otimes_{\Pi'}^{s''} X_2) = Y$  and  $s' = s'''$ .